EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
S1	5286	content with distribution with network	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2007/07/07 17:21
S2	582168	mirror$	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2007/07/07 17:21
S3	383700	partition$	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2007/07/07 17:21
S4	47687	load with balanc$	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2007/07/07 17:21
S5	3687	S2 same S3	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2007/07/07 17:21
S6	8	S1 and S5	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2007/07/07 17:21
S7	959	S4 same S3	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2007/07/07 17:21
S8	1	S6 and S7	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2007/07/07 17:22
S9	13604544	@ad<"20030625"	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2007/07/07 17:22
S10	21701	microsoft.as.	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2007/07/07 17:22
S11	13596175	S9 not S10	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2007/07/07 17:22

EAST Search History

S12	5487	S2 and S4	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2007/07/07 17:22
S13	3675	S11 and S12	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2007/07/07 17:22
S14	100	S1 and S13	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2007/07/07 17:23
S15	57	S3 and S14	US-PGPUB; USPAT; EPO; JPO; IBM_TDB	OR	ON	2007/07/07 17:23

PALM INTRANET

Day : Saturday
Date: 7/7/2007
Time: 19:40:37

Inventor Information for 10/606626

Inventor Name	City	State/Country
DALIA, APURVA F.	KIRKLAND	WASHINGTON
HARRY, CRAIG ALLEN	WAKE FOREST	NORTH CAROLINA
DANI, NISHANT	ISSAQUAH	WASHINGTON
JIANG, SHAN	BELLEVUE	WASHINGTON
THOMPSON, BRAD DEAN	WOODINVILLE	WASHINGTON
BARROWS, BRADLEY J.	REDMOND	WASHINGTON
SHUTT, DAVID R.	SEATTLE	WASHINGTON

Appln Info | Contents | Petition Info | Atty/Agent Info | Continuity/Reexam | Foreign D

Search Another: Application # Search **or Patent#** Search

PCT / / Search **or PG PUBS #** Search

Attorney Docket # Search

Bar Code # Search

To go back use Back button on your browser toolbar.

Back to PALM | ASSIGNMENT | OASIS | Home page

☐ **Search Session History**

Edit an existing query or compose a new query in the Search Query Display.

Select a search number (#) to:

- Add a query to the Search Query Display
- Combine search queries using AND, OR, or NOT
- Delete a search
- Run a search

Sat, 7 Jul 2007, 7:56:44 PM EST

Search Query Display

Recent Search Queries

#1	(((mirroring<in>metadata) <and> (load balancing<in>metadata))<and> (content distribution network<in>metadata)) <and> (pyr >= 1989 <and> pyr <= 2003)
#2	(((mirroring<in>metadata) <and> (load balancing<in>metadata))) <and> (pyr >= 1997 <and> pyr <= 2003)

Indexed by Inspec

Published before June 2003
Terms used: **load balancing mirroring cdn**

Found **8,627** of **143,151**

Sort results by relevance
Display results expanded form
Save results to a Binder
Search Tips
Open results in a new window

Try an Advanced Search
Try this search in The ACM Guide

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next
Best 200 shown
Relevance scale

1 Application level performance: On the use and performance of content distribution networks
Balachander Krishnamurthy, Craig Wills, Yin Zhang
November 2001 **Proceedings of the 1st ACM SIGCOMM Workshop on Internet Measurement IMW '01**
Publisher: ACM Press
Full text available: pdf(2.51 MB) Additional Information: full citation, abstract, references, citings, index terms

Content distribution networks (CDNs) are a mechanism to deliver content to end users on behalf of origin Web sites. Content distribution offloads work from origin servers by serving some or all of the contents of Web pages. We found an order of magnitude increase in the number and percentage of popular origin sites using CDNs between November 1999 and December 2000.In this paper we discuss how CDNs are commonly used on the Web and define a methodology to study how well they perform. A performanc ...

2 Network behavior: The effectiveness of request redirection on CDN robustness
Limin Wang, Vivek Pai, Larry Peterson
December 2002 **ACM SIGOPS Operating Systems Review**, Volume 36 Issue SI
Publisher: ACM Press
Full text available: pdf(1.86 MB) Additional Information: full citation, abstract, references, cited by, index terms

It is becoming increasingly common to construct network services using redundant resources geographically distributed across the Internet. Content Distribution Networks are a prime example. Such systems distribute client requests to an appropriate server based on a variety of factors---e.g., server load, network proximity, cache locality--in an effort to reduce response time and increase the system capacity under load. This paper explores the design space of strategies employed to redirect reque ...

3 Consistency and replication: Modeling redirection in geographically diverse server sets
Lisa Amini, Anees Shaikh, Henning Schulzrinne
May 2003 **Proceedings of the 12th international conference on World Wide Web WWW '03**
Publisher: ACM Press

Full text available: pdf(362.44 KB) Additional Information: full citation, abstract, references, citings, index terms

Internet server selection mechanisms attempt to optimize, subject to a variety of constraints, the distribution of client requests to a geographically and topologically diverse pool of servers. Research on server selection has thus far focused primarily on techniques for choosing a server from a group administered by single entity, like a content distribution network provider. In a federated, multi-provider computing system, however, selection must occur over distributed server sets deployed by ...

Keywords: content distribution network (CDN), performance, server selection, web traffic redirection

4 The state of the art in locally distributed Web-server systems

Valeria Cardellini, Emiliano Casalicchio, Michele Colajanni, Philip S. Yu
June 2002 **ACM Computing Surveys (CSUR)**, Volume 34 Issue 2
Publisher: ACM Press

Full text available: pdf(1.41 MB) Additional Information: full citation, abstract, references, citings, index terms

The overall increase in traffic on the World Wide Web is augmenting user-perceived response times from popular Web sites, especially in conjunction with special events. System platforms that do not replicate information content cannot provide the needed scalability to handle large traffic volumes and to match rapid and dramatic changes in the number of clients. The need to improve the performance of Web-based services has produced a variety of novel content delivery architectures. This article w ...

Keywords: Client/server, World Wide Web, cluster-based architectures, dispatching algorithms, distributed systems, load balancing, routing mechanisms

5 Dynamic services and analysis: Engineering and hosting adaptive freshness-sensitive web applications on data centers

Wen-Syan Li, Oliver Po, Wang-Pin Hsiung, K. Selçuk Candan, Divyakant Agrawal
May 2003 **Proceedings of the 12th international conference on World Wide Web WWW '03**
Publisher: ACM Press

Full text available: pdf(10.31 MB) Additional Information: full citation, abstract, references, citings, index terms

Wide-area database replication technologies and the availability of content delivery networks allow Web applications to be hosted and served from powerful data centers. This form of application support requires a complete Web application suite to be distributed along with the database replicas. A major advantage of this approach is that dynamic content is served from locations closer to users, leading into reduced network latency and fast response times. However, this is achieved at the expense ...

Keywords: database-driven web applications, dynamic content, freshness, response time, net-work latency, web acceleration

6 Enabling dynamic content caching for database-driven web sites

K. Selçuk Candan, Wen-Syan Li, Qiong Luo, Wang-Pin Hsiung, Divyakant Agrawal
May 2001 **ACM SIGMOD Record , Proceedings of the 2001 ACM SIGMOD international conference on Management of data SIGMOD '01**, Volume 30 Issue 2
Publisher: ACM Press

Full text available: pdf(319.67 KB) Additional Information: full citation, abstract, references, citings, index terms

Web performance is a key differentiation among content providers. Snafus and slowdowns at major web sites demonstrate the difficulty that companies face trying to scale to a large amount of web traffic. One solution to this problem is to store web content at server-side and edge-caches for fast delivery to the end users. However, for many e-commerce sites, web pages are created dynamically based on the current state of business processes, represented in application servers and *databases*
Keywords: JDBC, application server, database driven web site, dynamic content caching, invalidation, web acceleration

7 Higher-order distributed objects

Henry Cejtin, Suresh Jagannathan, Richard Kelsey
September 1995 **ACM Transactions on Programming Languages and Systems (TOPLAS)**, Volume 17 Issue 5
Publisher: ACM Press

Full text available: pdf(2.33 MB) Additional Information: full citation, abstract, references, citings, index terms, review

We describe a distributed implementation of Scheme that permits efficient transmission of higher-order objects such as closures and continuations. The integration of distributed communication facilities within a higher-order programming language engenders a number of new abstractions and paradigms for distributed computing. Among these are user-specified load-balancing and migration policies for threads, incrementally linked distributed computations, and parameterized client-server applicat ...

Keywords: concurrency, continuations, higher-order languages, message-passing

8 Hot mirroring: a method of hiding parity update penalty and degradation during rebuilds for RAID5

Kazuhiko Mogi, Masaru Kitsuregawa
June 1996 **ACM SIGMOD Record , Proceedings of the 1996 ACM SIGMOD international conference on Management of data SIGMOD '96**, Volume 25 Issue 2
Publisher: ACM Press

Full text available: pdf(1.37 MB) Additional Information: full citation, abstract, references, citings, index terms

This paper proposes a storage management scheme for disk arrays, named hot mirroring. In this scheme, storage space is partitioned into two regions. One is the mirrored region, which is characterized by high performance and low storage efficiency. The other is the RAID5 region, which is characterized by low performance and high storage efficiency. Hot data blocks are stored in the former area, while cold blocks are stored in the latter. In addition, mirrored pairs and RAID5 stripes are orthogona ...

9 Data partitioning and load balancing in parallel disk systems

Peter Scheuermann, Gerhard Weikum, Peter Zabback
February 1998 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 7 Issue 1
Publisher: Springer-Verlag New York, Inc.

Full text available: pdf(310.27 KB) Additional Information: full citation, abstract, citings, index terms

Parallel disk systems provide opportunities for exploiting I/O parallelism in two possible ways, namely via inter-request and intra-request parallelism. In this paper, we discuss the main issues in performance tuning of such systems, namely striping and load balancing, and show their relationship to response time and throughput. We outline the

main components of an intelligent, self-reliant file system that aims to optimize striping by taking into account the requirements of the applications, an ...

Keywords: Data allocation, Disk cooling, File striping, Load balancing, Parallel disk systems, Performance tuning

10 Aging through cascaded caches: performance issues in the distribution of web content

Edith Cohen, Haim Kaplan

August 2001 **ACM SIGCOMM Computer Communication Review , Proceedings of the 2001 conference on Applications, technologies, architectures, and protocols for computer communications SIGCOMM '01**, Volume 31 Issue 4

Publisher: ACM Press

Full text available: pdf(327.13 KB) Additional Information: full citation, abstract, references, citings, index terms

The Web is a distributed system, where data is stored and disseminated from both *origin* servers and *caches*. Origin servers provide the most up-to-date copy whereas caches store and serve copies that had been cached for a while. Origin servers do not maintain per-client state, and weak-consistency of cached copies is maintained by the origin server attaching to each copy an expiration time. Typically, the lifetime-duration of an object is fixed, and as a result, a copy fetched direc ...

11 Load distribution among replicated Web servers: a QoS-based approach

Marco Conti, Enrico Gregori, Fabio Panzieri

March 2000 **ACM SIGMETRICS Performance Evaluation Review**, Volume 27 Issue 4

Publisher: ACM Press

Full text available: pdf(695.59 KB) Additional Information: full citation, abstract, index terms

A dominant factor for the success of an Internet based Web service is the Quality of Service (QoS) perceived by its users. The principal QoS attributes these users perceive include those related to the service "responsiveness", i.e. the service availability and timeliness. In this paper, we argue that QoS can be provided by distributing the processing load among replicated Web servers, and that these servers can be geographically distributed across the Internet. In this context, we discuss strat ...

Keywords: QoS, Web server, load distribution

12 Using rotational mirrored declustering for replica placement in a disk-array-based video server

Ming-Syan Chen, Hui-I Hsiao, Chung-Sheng Li, Philip S. Yu

January 1995 **Proceedings of the third ACM international conference on Multimedia MULTIMEDIA '95**

Publisher: ACM Press

Full text available: htm(46.60 KB) Additional Information: full citation, references, citings, index terms

13 Cluster I/O with River: making the fast case common

Remzi H. Arpaci-Dusseau, Eric Anderson, Noah Treuhaft, David E. Culler, Joseph M. Hellerstein, David Patterson, Kathy Yelick

May 1999 **Proceedings of the sixth workshop on I/O in parallel and distributed systems IOPADS '99**

Publisher: ACM Press

Full text available: pdf(1.20 MB) Additional Information: full citation, references, citings, index terms

14 On robust transaction routing and load sharing

Philip S. Yu, Avraham Leff, Yann-Hang Lee
September 1991 **ACM Transactions on Database Systems (TODS)**, Volume 16 Issue 3
Publisher: ACM Press

Full text available: pdf(2.49 MB) Additional Information: full citation, abstract, references, citings, index terms, review

In this paper we examine the issue of robust transaction routing in a locally distributed database environment where transaction characteristics such as reference locality imply that certain processing systems can be identified as being more suitable than others for a given transaction class. A response time based routing strategy can strike a balance between indiscriminate sharing of the load and routing based only on transaction affinity. Since response time estimates depend on workload a ...

Keywords: distributed database, load balancing, performance analysis, transaction routing

15 RAID: high-performance, reliable secondary storage

Peter M. Chen, Edward K. Lee, Garth A. Gibson, Randy H. Katz, David A. Patterson
June 1994 **ACM Computing Surveys (CSUR)**, Volume 26 Issue 2
Publisher: ACM Press

Full text available: pdf(3.60 MB) Additional Information: full citation, abstract, references, citings, index terms, review

Disk arrays were proposed in the 1980s as a way to use parallelism between multiple disks to improve aggregate I/O performance. Today they appear in the product lines of most major computer manufacturers. This article gives a comprehensive overview of disk arrays and provides a framework in which to organize current and future work. First, the article introduces disk technology and reviews the driving forces that have popularized disk arrays: performance and reliability. It discusses the tw ...

Keywords: RAID, disk array, parallel I/O, redundancy, storage, striping

16 Run-time adaptation in river

Remzi H. Arpaci-Dusseau
February 2003 **ACM Transactions on Computer Systems (TOCS)**, Volume 21 Issue 1
Publisher: ACM Press

Full text available: pdf(849.04 KB) Additional Information: full citation, abstract, references, index terms

We present the design, implementation, and evaluation of run-time adaptation within the River dataflow programming environment. The goal of the River system is to provide adaptive mechanisms that allow database query-processing applications to cope with performance variations that are common in cluster platforms. We describe the system and its basic mechanisms, and carefully evaluate those mechanisms and their effectiveness. In our analysis, we answer four previously unanswered and important que ...

Keywords: Performance availability, clusters, parallel I/O, performance faults, robust performance, run-time adaptation

17 A comparison of high-availability media recovery techniques

George Copeland, Tom Keller
June 1989 **ACM SIGMOD Record , Proceedings of the 1989 ACM SIGMOD international conference on Management of data SIGMOD '89**, Volume 18 Issue 2
Publisher: ACM Press

Full text available: pdf(1.32 MB) Additional Information: full citation, abstract, references, citings, index terms

We compare two high-availability techniques for recovery from media failures in database systems. Both techniques achieve high availability by having two copies of all data and indexes, so that recovery is immediate. "Mirrored declustering" spreads two copies of each relation across two identical sets of disks. "Interleaved declustering" spreads two copies of each relation across one set of disks while keeping both copies of each tuple on separate disks. Both ...

18 Asynchronous scheduling of redundant disk arrays

Peter Sanders
July 2000 **Proceedings of the twelfth annual ACM symposium on Parallel algorithms and architectures SPAA '00**
Publisher: ACM Press

Full text available: pdf(161.35 KB) Additional Information: full citation, abstract, references, citings, index terms

Random redundant allocation of data to parallel disk arrays can be exploited to achieve low access delays. New algorithms are proposed which improve the previously known shortest queue algorithm by systematically exploiting that scheduling decisions can be deferred until a block access is actually started on a disk. These algorithms are also generalized for coding schemes with low redundancy. Using extensive experiments, practically important quantities are measured which have so far eluded ...

19 Cluster-based scalable network services

Armando Fox, Steven D. Gribble, Yatin Chawathe, Eric A. Brewer, Paul Gauthier
October 1997 **ACM SIGOPS Operating Systems Review , Proceedings of the sixteenth ACM symposium on Operating systems principles SOSP '97**, Volume 31 Issue 5
Publisher: ACM Press
Full text available: pdf(2.42 MB) Additional Information: full citation, references, citings, index terms

20 Compilers I: Affinity-based cluster assignment for unrolled loops

Gayathri Krishnamurthy, Elana D. Granston, Eric J. Stotzer
June 2002 **Proceedings of the 16th international conference on Supercomputing ICS '02**
Publisher: ACM Press

Full text available: pdf(633.13 KB) Additional Information: full citation, abstract, references, citings, index terms

To compete performance-wise, modern VLIW processors must have fast clock rates and high instruction-level parallelism (ILP). Partitioning resources (functional units and registers) into clusters allows the processor to be clocked faster, but operand transfers across clusters can easily become a bottleneck. Increasing the number of functional units increases the potential ILP, but only helps if the functional units can be kept busy.To support these features, optimizations such as loop unrolling m ...

Keywords: VLIW architectures, affinity-based clustering (ABC) algorithms, cluster assignment, homogeneous clusters, loop optimizations, loop scheduling, loop unrolling, partitioned register files, software pipelining

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next